SearchNetworking.com

The ESSENTIAL GUIDE for Upgrading your Network

CHAPTER 1

Networking Evolution and Roadmap

By David Greenfield

The network is on the move yet again. Running voice, videoconferencing and video surveillance across the network that also handles email, file transfers and critical business transactions is enabling organizations to reduce overall capital expenditures and improve their agility.

But none of that comes for free. A predictable infrastructure is needed for these new applications to work. Packet loss, delay and jitter must be rigorously inhibited if a pin drop is to be heard across a phone call. Those requirements are only going to become stiffer as application load and user expectations grow.

Availability is crucial. CEOs aren't particularly fond of having to reboot their phones. They don't want their surveillance systems to fail. And they most certainly aren't interested in hearing why their point-of-sale systems are offline or why a manufacturing plant has had to shut down. So before IT professionals can expect to deploy business-critical applications across their networks, they need to be sure that the network is business-critical ready. It's not just about having a huge pipe; it's about delivering a complete facility that supplies applications consistently, reliably and securely.

The art of impact assessments

Ensuring a properly designed infrastructure is the first step toward an effective deployment, be it of voice, video or any other application. Ultimately, though, any thorough impact assessment is more than just answering a checklist of questions. It relies on a deep understanding of the business problem that must be solved. Only then can the underlying technical requirements be determined.

Do regulatory requirements around data preservation need to be met? Are assurances around customer delivery metrics being promised, and must they be met? Or, even more simply, is that telephone being deployed sitting on the CEO's desk or the janitor's?

Answering questions such as these is fundamental to organizations' leveraging their infrastructure investment. If a business does not answer these questions, IT will find it very difficult to gauge the status of the current network, prioritize upgrades and evaluate whether improvements are meeting

IT and business needs. An impact assessment detailing business requirements and outlining a roadmap to transition to the architecture that will fulfil those requirements in several main areas is essential to success.

Physical plant

Complex, modern networks cannot run well without a solid underlying physical plant. Beefing up wiring, switches and routers can make a big difference in overall network performance.

Be certain to conduct a physical layer audit to ensure that all installed cabling is up to par. It's amazing how many stories of transient VoIP problems resolve into a question of bad cabling, a faulty splitter or a loose coupling. Take the time to be sure that the cabling plant meets specifications, and correct any errors before deploying those uber-applications.

While mobility is becoming a way of life for many users, most typical LANs are based on good old-fashioned Cat-5 or Cat-6 cabling. It's entirely probable that a network upgrade in three to five years may use wireless for more than guest Internet access; but today, the challenges of delivering quality of service, roaming between access points, security, voice quality and numerous other issues make wireless ill suited as the basis of many corporate networks. Networking professionals should ensure that wireless access is reliable and consistent in the places where it is deployed, but in most instances the wired network will be carrying the bulk of business-critical traffic.

Bandwidth: How much is enough?

Without the right amount of bandwidth in the network, real-time applications will grind to a halt. Finding what's right can be a challenge. One easy answer in the LAN environment is to roll out switched, 100 Mbps Ethernet to the desktop. Upgrading to 1 Gbps to the desktop is another option, but unless an organization moves huge files, such as engineering designs, most desktops won't be able to take advantage of the additional bandwidth.

Without the right amount of bandwidth in the network, real-time applications will grind to a halt.

The bigger challenge is providing enough bandwidth at key network junction points in the network, as well as over the WAN. Typically, there

is a mismatch between incoming demand and the outbound capacity. That can create bottlenecks where any number of links aggregate together, such as the switch at the front end of a server, or where the LAN meets the WAN.

Ensuring optimum performance on a minimal budget requires a good understanding ... of the applications running over the network.

For these instances, ensuring optimum performance on a minimal budget requires a good understanding of the behavior of the applications running over the network. Understanding the application's latency, jitter and packet loss requirements, as well as its traffic patterns—such as whether flows are peer-to-peer, involve a server, or both—is critical for leveraging the capacity in your network topology.

Ultimately, 10 Gbps links may clear away any performance problems within the LAN, most notably to servers, especially those using virtualization. Today, however, 10 Gbps is still cost prohibitive for many organizations. For them, a more realistic option is sticking with 1 Gbps links and then distributing traffic across multiple servers using new technology such as application delivery controllers and, ultimately, grid or virtualized architectures.

WAN bandwidth must also be carefully considered. Businesses are global and mobile, making bandwidth use between locations skyrocket, and trends toward centralized applications and data center consolidation have created more of those congestion points where the LAN meets the WAN. Enterprises can work with carriers to negotiate the most cost-effective and appropriate WAN links. They can also implement WAN optimization and acceleration to use the WAN as efficiently as possible.

Predictable delivery

Invariably, someone in the organization will want to save a bit of money on the network upgrade by challenging the need for new switches with quality of service (QoS) for prioritizing traffic. He or she will point out that, even when VoIP is running on the network, the amount of bandwidth a VoIP session uses with a bulky wideband CODEC is a fraction of the speed of even a 10 Mbps link, let alone a 100 Mbps link.

Insist on QoS. While it may be true that voice doesn't require QoS under "normal" conditions, the network should not be designed for normal conditions. You must plan for peak conditions, such as a virus spreading across the network or a user deciding to back up a 100 GB file across the network. For those instances, the network staff must ensure that VoIP will take higher priority than Internet browsing.

While an issue for LANs, QoS becomes an even bigger issue in any location where there is a speed mismatch, such as on the access link from a customer premises to the Internet or the corporate WAN. With smaller bandwidth capacities, access links are notorious for places where delay-sensitive applications falter. Implemented properly, QoS prevents applications from hogging those links and ensures that priority traffic gets through.

Bandwidth consistency also must be considered in the context of network devices. The delay through a router can interrupt applications as well, and routers offer many queuing mechanisms for reducing transit times. Although routers may perform fine under normal conditions, however, it's the unusual or extreme conditions that IT needs to consider. When routers recalculate their routing tables in corporate WANs, for example, routing delays can result. Examining these "boundary conditions" will go a long way toward preparing your routing networking infrastructure for real-time applications.

Reliability

Reliability is a cornerstone of an enterprise-class network. In a network running VoIP, availability and reliability are complicated by having to protect the logical as well as physical facilities. Good design suggests running voice in its own virtual LAN (VLAN) to protect voice calls from the vagaries of data applications running on the rest of the network. However, doing so also allows for network architectures that may build redundancy in the switching fabric or routing core but still be unable to complete a voice call.

With smaller bandwidth capacities, access links are notorious for places where delay-sensitive applications falter.

Ensuring reliability starts with the physical network. Classic tiered-switching design provides an inherent

measure of failover. Should the distribution switch connecting a workgroup's access switch to the network's core switch fail, a second distribution switch is typically designed into the overall network plan to handle the load from the access switch. But such an approach consumes additional switch ports, the costs of which must be considered.

Simply segmenting voice traffic cannot protect a voice system nor can simply implementing a firewall secure the network from external attackers.

Router availability is a similar issue. Numerous options exist for individual router survivability and operation in the event of a failure. Alternatively, the networking team may consider redundant router designs as a more effective route.

Power supply and environmental factors within the writing closet must also be carefully considered. All devices will have to be backed up by sufficient power to meet the organization's needs during an outage. Again, the duration of power depends on the organization's objectives, but two to four hours is typical within the industry. If switches are to offer inline power, such as Power over Ethernet (PoE), the power requirements will be higher. This, in turn, means ensuring sufficient power delivery and cooling requirements for the wiring closet. It also means considering how to power end nodes during an outage.

As for logical failover, VLAN assignments must be done per node, not per application. While telephones may be easily assigned to a voice VLAN, softphones and unified communications (UC) clients are another matter. They are typically assigned to a data VLAN. Enterprises then need to route between VLANs if phones, softphones and UC clients are to communicate with one another.

Security

VLANs provide a small measure of security, but simply segmenting voice traffic cannot protect a voice system nor can simply implementing a firewall secure the network from external attackers. The distributed nature of modern networks, combined with the increased focus on applications and business data, makes network security less about locking down a network perimeter and more about controlling user access and behavior and ensur-

ing data safety in a dynamic and mobile environment.

Today, more attacks come from within the network than outside it. A survey last year by Deloitte showed that 83% of IT executives at technology, media and telecommunications companies were concerned about "employee misconduct involving information systems."

Protecting against these users requires rethinking today's security architecture. IT must assume that desktops have been compromised. Central to that protection architecture is the use of tools like network access control (NAC) and endpoint security products. Using these tools, centralized security systems can ensure that only authorized individuals are granted access to the network using approved devices and that those devices are using protocols, applications and processes defined by the enterprise.

IT can also implement data protection restrictions in the form of digital rights management (DRM). The combination of the endpoint security and DRM allows organizations to restrict access when data is editable on a user's machine and limit how it may be distributed across the network.

These measures are implemented in addition to security practices and systems that already exist in most organizations. Intrusion detection and prevention capabilities, for example, are needed for post-admission control, as are use of access control lists (ACLs) and internal firewalls to limit access to sensitive parts of the business.

Central to a data protection architecture is the use of tools like network access control (NAC) and endpoint security products.

Many companies have implemented Secure Sockets Layer (SSL) VPN for secure remote access for telecommuting and mobile workers.

Thought must also be given to how calls are placed to and from the PSTN. Today, that is normally done by purchasing local gateways but, increasingly, organizations are considering IP-based trunks supplied by an external service provider. This eliminates the cost of gateways while enabling calls to remain on the IP network end-to-end, improving call quality. At the same time, however, it exposes the network to external IP sessions and requires additional VoIP security. ■

David Greenfield is principal of STAnalytics, a global technology marketing consultancy where he advises enterprises on emerging technologies. He has spent the past 20 years analyzing network communications and most recently was the editor of *Network Computing*. His work has appeared in other leading technology publications, such as *PC Magazine*, *IT Architect*, *Data Communications* and *Red Herring*, and he has consulted to and assisted Fortune 500 enterprises in their technology acquisitions.